

AD-A119 741

BOLT BERANEK AND NEWMAN INC CAMBRIDGE MA
ON INTERPRETING SEMANTIC NETWORK FORMALISMS.(U)
SEP 82 O J ISRAEL
BBN-5117

F/6 5/7

N00014-77-C-0371

NL

UNCLASSIFIED

[OF]

AD &
09741

END

DATE

FILED

1-82

DTIC

Bolt Beranek and Newman Inc.

(12)

BBN

AD A119741

Report No. 5117

On Interpreting Semantic Network Formalisms

David J. Israel

September 1982

Prepared for:
The Office of Naval Research

DTIC FILE COPY

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

82

003

Report No. 5117

ON INTERPRETING SEMANTIC NETWORK FORMALISMS*

David J. Israel

September 1982

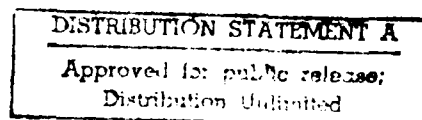


Prepared by:

Bolt Beranek and Newman Inc.
10 Moulton Street
Cambridge, MA 02238

Prepared for:

The Office of Naval Research



This research was supported by the Office of Naval Research under Contract No. N00014-77-C-0371. The views and conclusions contained in this document are those of the author and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Office of Naval Research or the U.S. Government.

*A version of this report is to appear in a special issue on Computational Linguistics of the International Journal of Computers and Mathematics.

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER BBN Report No. 5117	2. GOVT ACCESSION NO. AD-A119 741	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) ON INTERPRETING SEMANTIC NETWORK FORMALISMS		5. TYPE OF REPORT & PERIOD COVERED Technical Report
7. AUTHOR(s) David J. Israel		6. PERFORMING ORG. REPORT NUMBER BBN Report No. 5117
9. PERFORMING ORGANIZATION NAME AND ADDRESS Bolt Beranek and Newman Inc. 10 Moulton Street Cambridge, MA 02238		8. CONTRACT OR GRANT NUMBER(s) N00014-77-C-0371
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Department of the Navy Arlington, VA 22217		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE September 1982
		13. NUMBER OF PAGES 49
		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Distribution of this document is unlimited. It may be released to the Clearinghouse, Dept. of Commerce, for sale to the general public.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Knowledge representations, semantics, semantic networks, representational formalisms.		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) In a recent paper, Reiter and Criscuolo remark "that (semantic) networks are notational variants of logical formulae is by now a truism in Artificial Intelligence circles". Shamelessly exploiting the foregoing quote as a pretext, I attempt to sketch adequate semantic accounts for at least two (kinds of) semantic network formalisms: one, based on the notion of inheritance; one, not. A crucial condition of adequacy to be		

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

20. Abstract (cont'd.)

satisfied is fidelity to some of the intuitions of the creators of the formalisms.

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Avail and/or	
Dist	Special
R	

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

TABLE OF CONTENTS

	Page
1. INTRODUCTION	3
2. SNePS	7
3. INHERITANCE	17
3.1 THE SEMANTICS OF INHERITANCE	20
4. TECHNICAL ADDENDUM: TAXONOMIES OF NATURAL KINDS	39
5. ACKNOWLEDGEMENTS	45

ABSTRACT

In a recent paper, Reiter and Criscuolo remark "that (semantic) networks are notational variants of logical formulae is by now a truism in Artificial Intelligence circles" [24]. Shamelessly exploiting the foregoing quote as a pretext, I attempt to sketch adequate semantic accounts for at least two (kinds of) semantic network formalisms: one, based on the notion of inheritance; one, not. A crucial condition of adequacy to be satisfied is fidelity to some of the intuitions of the creators of the formalisms.

DOCTRINAL PREAMBLE

One often hears that modal (or some other) logic is pointless because it can be translated into some simpler language in a first-order way. Take no notice of such arguments. There is no weight to the claim that the original system must therefore be replaced by the new one. What is essential is to single out important concepts and to investigate their properties. The fact that the real numbers can be defined in terms of sets is no argument for being interested in arbitrary sets. One must look among the sets for the significant ones and cannot be censured if one finds the intrinsic properties of the reals more interesting than any of their formulations in set theory. Of course if we can argue that set theory provides other significant concepts, then we may find some reason for going beyond the real numbers (and it is not hard to find the reasons!). But clearly this discussion cannot proceed on purely formal grounds alone. (From [26]).

AMEN!

1. INTRODUCTION

In a recent paper, Ray Reiter and Giovanni Criscuolo remark that "the fact that networks are notational variants of logical formulae is by now a truism in Artificial Intelligence circles" [24]. Let us put aside the empirical, sociological claim - about which I am more than willing to defer to Messrs. Reiter and Criscuolo. Let us look rather at the content of the truism itself.

When we do so, we notice a certain ambiguity. Perhaps, Reiter and Criscuolo are supposing that there is at least one logical language for whose formulae the "sentential analogues" of each and every semantic network formalism are notational variants. Or, switching the quantifiers, perhaps the claim is that for every semantic network formalism there is at least one logical language of which it is a notational variant. Now these are two very different claims. (One could easily imagine a third, stronger claim, being presupposed, viz., that all semantic network schemes are notational variants of, in particular, the language and logic of classical, first-order quantification theory. I will, in what follows, address myself to this strongest, most specific, view as well.)

I am not going to speculate on which of the three views Reiter and Criscuolo really had in mind. Indeed, the quote from Reiter and Criscuolo is the merest pretext for the present paper. I intend to sketch quite different semantic accounts for two different (kinds of) semantic network formalisms - one organized around the notion of inheritance, the other, not; the latter, keyed to a certain family of intensional contexts, the former, not. Most important, the semantic accounts will differ and, with

luck, the differences will reflect differences in the central intuitions of the semantic net theorists concerned. The aim is, simply, to try to take semantic net theorists at their word; and to show that, in doing so, one must range fairly far and wide beyond the confines of standard first-order logic. I am not going to argue for adopting any particular semantic net formalism; nor am I going to examine any of the intuitions motivating the work to be discussed.

To make clear my intentions, let me note two ways in which one might fail to take the "semantic(s)" in "semantic network" seriously. The first is represented by the work of Quillian, Collins, et al., work which originated the semantic network tradition. By my lights, the structures described in this work are not intended to be languages; rather they are part of a theory or model of a certain range of memory-related psychological phenomena. The nodes in the network might be (associated with) words, or even sets of sentences, but the accounts of the network are not semantic accounts; they do not constitute a semantic theory. There is no attempt to account for the meaningfulness, or describe the meaning, of the nodes associated with complex "syntactic" entities in terms of the meanings of nodes for "syntactic" simples (plus a fixed set of constructors). This, of course, is no criticism of the work; nor am I suggesting that the researchers in question were confused about the present point. So much for the first way of not taking the "semantics" of semantic network formalisms seriously.

As a second way of not taking the semantics of semantic net formalisms seriously, I have in mind the following doctrine: the only way to interpret semantic network formalisms semantically, no matter what the semantic network theorists may say, is to treat them as notational variants of standard first-order

languages, with their standard "Tarski-style" semantics. This might best be described as a way of not taking semantic network theorists seriously. It is against this view, in particular, that I mean to deploy the quotation from Dana Scott with which I started. Of course, for all I know, no one holds this view; in which case, I am arguing only against phantoms of my own fevered imagination. So much the better. I, at any rate, want to try to take some semantic net theorists at their very word. Let the appropriate semantic account fall out however it may.

Bolt Beranek and Newman Inc.

Report No. 5117

2. SNePS

I want first to discuss a system in which the notion of inheritance plays no (special) role. It would also be nice to have a case in which the theorists are both explicit about their motivating intuitions and diligent in presenting sufficient detail on which to hang semantic speculation. Sad to say, this narrows the field down quite a bit. Still, there are at least two choices; the (atypically nameless) system of Schubert, Cercone and Goebel [25] and Shapiro's SNePS [27, 28]. I have decided to examine the latter; and this, for two reasons. First, Shapiro (and Shapiro & Maida [30, 31]) presents a fairly explicit "philosophy of semantic networks", as well as an enormous body of detailed description of the workings of the system. I think the philosophy is widely shared and trying to account for the considerations operative in its formulation raises interesting problems. Second, Schubert, et al. are too explicit for my purposes. It's quite clear that they view their formalism as a Montague-style type-theoretic, intensional system. It would, I think, be illuminating to work out in detail a semantic account for the system described in [25]; but there can anyway be no doubt about the "logical space" within which we'd find ourselves. This is not the case, or so I shall claim, for SNePS. A semantic account appropriate to it will force us to wander into largely, though not completely unexplored, territory. (And this, I also claim, is no argument against SNePS. Remember the wisdom of the great Scott.)

Shapiro has argued that we should impose the following conditions on semantic networks:

1. each node represents a unique concept

2. each concept represented in the network is represented by a node
3. each concept represented in the network is represented by a unique node (the Uniqueness Principle)
4. arcs represent non-conceptual (logical ? - D.I.) binary relations between nodes. [30, 31]

In what follows I shall take Shapiro to be simply describing SNePS, thus ignoring his arguments for imposing these conditions on all semantic networks. I shall also (largely) ignore the fourth condition and, for that matter, a fifth: "the knowledge represented about each concept is represented by the structure of the entire network connected to the node representing that concept".

On the basis of these conditions, Shapiro (and Shapiro and Maida) contend that "all nodes of a semantic network represent only intensions" [30]. "Nodes of a semantic network represent unique intensional concepts" [31]. Again, I shall take this as describing SNePS, not as prescribing for all semantic networks.

Now what are we to make of this? Raising the question this way raises the issue as to whether my intentions toward SNePS are honorable. I hope they are. Appearances to the contrary notwithstanding, I am not singling out a few lines, ripped out of context, for malicious attention. First, no malice is intended and I trust no harm is done. Second, and more important, my reconstruction of SNePS attempts to embrace a large number of the claims and arguments in the texts. I will not attempt to support my claim in this respect by citation. I hope that any one who has read the papers will agree that I have presented at least one way of construing them - not the only way, and perhaps not the way favored by the authors. As for those who have not read the

material, I fear they shall have to take me at my (immediately preceding) word.

Let's remind ourselves of the project. We are to find a language-cum-semantics which can reasonably be taken to be that formal system, formulae of which correspond to the sentential pieces of SNePS. What, for instance, does the logical vocabulary of our target language consist of? Here we get help from [28, 29]. For our purposes, what's crucial is that the logical constants mentioned are generalizations of the familiar truth-functional connectives and quantifiers. But with respect to the latter, over what do bound variables range; what kinds of things are assigned to the variables?

We might as well start at the beginning and specify how atomic predication formulae are to be understood. The standard way of handling intensions in contemporary logic is to treat them as (represented by) functions from an index set of contexts or possible worlds into some set-theoretic construct on the domain of the model structure. So, to take an important instance, properties - the intensions associated with monadic predicates - are explicated as functions from the index set into subsets of the domain of possible individuals. We might, then, try imagining a model structure consisting of a domain D of possible entities, a non-empty set I (the index set or the set of possible worlds), and (optionally) a distinguished element w of I to represent the real world. Now define an individual concept in such a model structure as a function ic from I into D . (Total or partial? The traditional answer has always been total; but it is

not clear what answer is appropriate to SNePS.)¹

What, then, can we say about individual terms: about individual variables, individual constants, and definite descriptions of individuals? Given what Shapiro says, it is hard to see how there is any alternative to a uniform intensional treatment. In specifying models for SNePS, all such terms (including, *nota bene*, individual variables) get interpreted by being associated with individual concepts (not individuals, not members of *D*). A model for SNePS will associate with each individual constant an individual concept (a member of the set of functions from *I* into *D*) and assignments relative to such a model will do the same for individual variables. This means that the modal language-cum-logic is not of the standard variety. I said above that I would be assuming the general framework of Kripke-Montague style model-theoretic accounts; however, neither Kripke nor Montague propose semantic accounts in which the individual variables get assigned individual concepts. Dana Scott's advice that one opt for just such a uniform treatment of all individual terms [26], which Shapiro seems to be following, has been followed by just about no one else except Aldo Bressan in [6].

Enough about individual terms; how shall we handle predicate letters? Remember: all nodes are intensional. Individual terms are associated with individual concepts, not with possible individuals; so we can not, in good conscience, assign (e.g.) to

¹As we shall see, there are reasons for thinking this effort to reconstrue SNePS-style intensions in terms of Kripke-Montague model-theoretic treatments of intensional contexts somewhat misguided. I shall suppress these reasons till the end of my discussion of SNePS. As for such model-theoretic treatments themselves, the classic sources are [15, 16, 20].

one-place predicate letters, functions from I to subsets of D . That is, we can't assign to one-place predicate letters sets of possible individuals. The obvious move might seem to be to associate sets of individual concepts with monadic predicate letters; but this does not render "predicate-nodes" truly intensional. The extension of a predicate at a world is no doubt a set of individual concepts; but what is its intension? Surely, it is (as both Bressan and Scott insist) a function from I into the power set of the set of individual concepts. Or what comes to the same thing, a function from individual concepts into propositions, where these are functions from I into $\{T, F\}$. (Shapiro is explicitly committed to "propositional nodes" - nodes for "concepts of the TRUE" and "concepts of the FALSE" [30, 31].) So, predication is an intensional functor.

By the way, this does make it a little hard to understand what Shapiro says about MEMBER and CLASS arcs, and the relation between these and ISA arcs [28]. I assume ISA links represent the predication functor; but this functor must be intensional - that is, at each world i , the truth value of " Fx " (x is an F) is not a function solely of the extension of " x " in i . In general, the relation between " Fx " and " x is a member of the set $\{y: Fy\}$ " is complicated. In particular, it is relatively straightforward only for extensional predicates, predicates which informally meet the following condition: they apply to a given ic x (at i) iff they apply to any d such that $d = x(i)$. The truth value at a world of a sentence predicating an extensional predicate to an ic does depend only on the extension of that ic at that world.²

²Needless to say, identity between elements of D is world-relative; the primitive notion is strict identity between ic 's - that is, co-extensiveness at all i . For simplicity, we assume totality.

Skipping lots of nasty details, we are now in a position to wave our hands, with some confidence, over the first order quantifiers. But note the "first-order" in the foregoing. Shapiro explicitly mentions the availability of higher-order constructs in SNePS. Thus, he says we can have nodes representing the second-order concepts of a property being extensional and of a property being intensional [30]. It's not clear how high we can go in this vein, and for reasons of space I herewith demur.³

Finally, there is the problem of propositional concepts. There are problems, in particular, about sentences embedded in functors such as "Necessarily..." or "S believes that..." About the first, Shapiro doesn't have much to say. Later, I shall suggest that the reason why he doesn't throws light on the fact that many of the choices we've made in giving a semantic account of SNePS seem ill-motivated by the texts.

The only intensional contexts Shapiro (and Shapiro and Maida) discusses are propositional attitude contexts, those involving verbs such as "know", "believe", etc. These contexts are treated, moreover, as relational; that is, "know" , "believe" (taken in their sentential complement mode) are treated, not as intensional operators, but as two-place predicates whose relata are individual concepts of subjects and either propositions or

³I should note, though, that Bressan's is an omega-type intensional system, as are some of the Montague logics.

propositional concepts.⁴ None of this tells us very much about how they would treat the standard modalities. Why no mention of these modalities, the hard-core intensional operators?

The answer to the mystery of the missing modalities is to be found, I think, in the intuitions behind the Principle of Uniqueness. The crucial motivation is the view that the "nodes [of a semantic network] represent the concepts and beliefs of a thinking, reasoning, language using being (e.g. a human)". [30]. Hence the centrality of the propositional attitudes. But these generate contexts which are not just intensional, but hyperintensional (and perhaps to the nth degree). (The slightly garish term is due to Cresswell [10].) Hyperintensional contexts are those in which substitution of logically equivalent sentences or strictly identical terms is not guaranteed to preserve truth. Thus, to take a particularly startling case, it is arguable that from the fact that S believes that P and Q it does not follow that S believes that Q and P. Again, from that the fact that S believes that 4 is even, it does not follow that S believes that the square of 2 is even. (Or that 4 is not odd.) Now Shapiro and Maida certainly seem to view belief contexts as hyperintensional and it is this which gives sense to the Principle of Uniqueness: no two distinct nodes (represent expressions which) are inter substitutable *salve veritate* in all the contexts the language can generate. That is, every distinction between nodes made purely syntactically by the

⁴As a reminder, a proposition - the intension of a sentence - is a function from I into {T,F}. A propositional concept would be an individual concept whose extension at a world is a proposition; it is a function from I into functions from I into {T,F}.

language is a semantically significant distinction - there is some context which semantically "separates" any two distinct expressions. Thus Shapiro: "No two distinct nodes represent truly equal concepts."⁵ Given this, it is clear enough why they would want to treat belief-contexts relationally; for it seems there can be no question of a "logic of belief". There would seem, that is, to be no laws governing the behavior of a proposed logical functor (operator) for belief. (Roughly speaking: if $[xB](P \ \& \ Q)$ didn't formally entail $[xB](Q \ \& \ P)$, would anything formally entail anything else?)⁶ And, given the very same, it is clear enough why so little is said about the standard modalities, which are merely (not hyper-) intensional.

There is no very happy account of hyperintensionality from within the model-theoretic framework (or elsewhere); although there are attempts. A move first proposed by Carnap in [7] is to specify a finer-grained notion than logical equivalence, that of intensional isomorphism, in terms of the construction trees of complex expressions and the intensions assigned the constituent expressions. For example, take two sentences which are purely truth-functional tautologies, hence which are true in the same (namely, all) possible worlds. They might, however, be formed out of different constituents and these might have different

⁵ [30] Note: it's one thing to require that no two primitive terms are co-intensional; it's quite another to argue that no two terms - primitive or not - are such. This latter, though, seems to be the Shapiro-Maida position.

⁶It is not necessary to hold this view of the hyperintensionality of propositional attitude contexts to dissuade one from the logical operator position; see [21, 22]. But it sure is sufficient.

intensions, etc. One might then make stipulations about intentional isomorphism being preserved under order-switching of operands of commutative operators, etc. This move may handle the case of sentential conjunction raised above; but it doesn't handle the other case(s).⁷

To return to the central point: if one is focussing on propositional attitude - allegedly hyperintensional - contexts, it can seem like a waste of time to introduce model-theoretic accounts of intensionality at all. Thus the air of desperation about the foregoing attempt to use such an account (albeit a non-standard one) to explicate a semantic net formalism that is focused on the propositional attitudes.

(More than) enough has been said, I hope, to support my claim that in taking SNePS seriously, in particular in attempting to present a semantic account which honors some of the intuitions of its creator(s), one is led into rather interesting, if slightly forbidding, logical terrain. Of course, enough has also been said to prove beyond a shadow of a doubt that there are significant open problems to be solved before a fully adequate account can be given; the major one being that of providing an

⁷Doubts have been raised about the efficacy of such a move, especially with respect to the iteration of propositional attitudes over different subjects. In particular, try Mates's matrix: Let D and D' be two intensionally isomorphic sentences. Then the following are also intensionally isomorphic: (a) Whoever believes (that) D, believes (that) D. (b) Whoever believes (that) D, believes (that) D'. But nobody wonders whether anybody doubts that whoever believes D believes D; but some philosopher may well wonder whether anybody doubts that whoever believes D believes D'. [18] There are other problems with Carnap's proposed treatment of belief and assertion; for a classic review, see [8]; and for a counterproposal, see [9].

account of propositional attitude contexts which is both formally impeccable and at least a little plausible.⁸

⁸There have, of course, been significant attempts in this direction; in the AI literature, the outstanding candidates are [19] and [23]. Needless to say, there is much, much more that needs to be said; sad to say, it will here go unsaid.

3. INHERITANCE

I want now to take a look at those semantic network formalisms in which the notion of inheritance is central. Here, too, the work is motivated by a certain "family" of intuitions. I would not want to try to ascribe these intuitions to any one in particular and perhaps no one believes all of them. I will state them baldly and without comment. (If you'd like, you may imagine me to be making up a position out of whole cloth and then, perversely, imposing upon myself the duty of making formal sense out it.)

1. The graph-theoretic nature of semantic networks counts for something above and beyond "ease" of formulating and implementing access and retrieval algorithms.
2. There is something in principle wrong about the way in which standard semantic accounts separate languages from "theories" expressed in them - not enough is fixed by the specification of a language itself. In some sense, different languages implicate different theories.
3. Somehow the central role in thought (and language) of kind terms, in particular of natural kind terms, must be captured - and to do this, one must take seriously the fact that natural kinds come in families on which hierarchical, taxonomic relations are defined.

These three "intuitions" coalesce to form a certain perspective on semantic nets.⁹ This perspective often carries along with it a commitment to some theory or other of prototypes

⁹For a more detailed attempt to take the intuition about the centrality of natural kinds seriously, see the Technical Addendum at the end of the present report.

and this commitment is often understood, mistakenly, as constituting an essential constraint on any semantic account of such a semantic net formalism. (For more on this, see [12].) In what follows, I shall completely ignore all issues about prototypes.¹⁰

Some preliminary comment on (2) is called for. It has often seemed that there was a systematic confusion of language and theory evident in the work on semantic networks. One was never given a specification of the "language" neat; rather, what one got were notations for particular sentences and one was supposed to be able to go on from there. But how is one to know how to go on, unless one knows, at the very least, what is fixed and structural, as against what is subject to variation by way (e.g.) of varying meaning assignments?¹¹ Despite (or because of) worries about this confusion, I think it worthwhile to attempt a rational reconstruction of the intuition that a partial theory of the world is "directly" embedded in the languages we use in thinking and speaking about that world.

The best way to shed a little light on the second of our central intuitions is by contrast. The standard mode of specifying a quantificational language (or: language scheme)

¹⁰As an aside, the semantic net formalism I have most clearly and fully in mind is KL-ONE, which likewise eschews prototypes. See [4, 5].

¹¹Remember the quotation from [30]: "A semantic network models the knowledge structure of a thinking, reasoning, language using being. In this case, nodes represent the concepts and beliefs such a being would have." This certainly sounds like an identification of a given network with a particular creature's view of its world.

includes the specification of a typically infinite set of typically infinite sets of predicate letters of all possible arities. To interpret such a language in the classical model-theoretic way, one assigns a set as the domain for the variables and an extension to each predicate letter. In the standard vein, these assignments are to subsets of the n -place Cartesian product of the domain of the variables (for n -place predicates).¹²

The predicates, individually, are syntactically unstructured and, collectively, are no more than members of various unordered sets. That is to say, from a semantic point of view, there are no constraints on the interpretations assigned to the members of any set of predicate letters beyond that imposed by the arity of the individual predicates. (So e.g., two-place predicates must be assigned subsets of the set of all ordered pairs of members of the domain, etc.) In fact, though, there is a kind of constraint: there are to be no semantic interdependencies among the extensions of the predicates. That the assignments to predicate letters be independent of one another mandates that the specification of the language cannot impose any relationships among the extensions of the predicates. Hence, the standard scheme imposes a requirement of logical independence among atomic sentences.

For instance, if one has two one-place predicate letters,

¹²A point of terminology: logicians divide on usage here. I tend to speak of one and the same language being susceptible to many different meaning assignments (interpretations). Others speak, rather, of different - applied - languages as instances of the same language scheme. Anything said in one mode can be translated into the other. I leave it to the reader to make the requisite transformations.

"P" and "Q", one would not (indeed, on many accounts, should not) assign to "Q" the complement of the set assigned to "P"; that's a job for the negation operator. One can, of course, specify a theory in the language which has " $(x)(Qx \leftrightarrow \neg Px)$ " as a theorem; but one can also formulate in the same language a theory according to which "Q" and "P" are (e.g.) coextensive.¹³

The picture sketched above should be familiar, but it does not seem to fit very well that body of work by researchers in Artificial Intelligence which focuses on the related notions of taxonomic structures and inheritance. I shall now present a mildly non-standard account in which those notions are indeed central.

3.1 THE SEMANTICS OF INHERITANCE

You may think of the sentences of the language(s) to be studied as looking a lot like the sentences of a standard first-order language, with a few wrinkles, of course. There will be a finite number of primitive predicates of not all possible

¹³We couldn't, in the foregoing, talk of the co-extensiveness of "Q" and "not-P"; at least not if we were restricting ourselves to the resources available in standard first-order languages. For complex predicates are not among those resources; so, in strictu sensu, there is no predicate available to play the role - solely in virtue of its structure or solely in virtue of the specification of the language and the standard semantic account - of the complement of another predicate. Rather, we must make do with (complex) open sentences, such as " $\neg Px$ ".

degrees.¹⁴ The crucial feature of the account resides in the requirement that integral to specifying the semantics of the language is the specification of an algebraic structure of properties (concepts, intensions¹⁵) by which the assignment of extensions to the predicates is constrained.

One more preliminary point: as noted just above, the elements of the algebraic structures are to be understood as properties. The "linguistic" representatives of these are, in the first instance, lambda-abstracts interpreted as singular terms denoting properties. These are then associated with monadic predicate letters of the language. When we get around to exploiting lambda-abstraction as a complex predicate-forming operator, predicates of the language will look a lot like the singular terms for properties with which they are correlated. This is unfortunate and could easily be remedied by choosing a different notation for the complex predicate forming operator. (But it won't be so remedied.) Occasional reminders of the distinction between singular terms for properties, on the one hand, and predicates, on the other, will be sprinkled about.

In the first instance, we shall limit ourselves to properties, properly so-called; that is, syntactically speaking,

¹⁴We will be able to generate non-primitive predicates of any arity by way of composition, e.g. by forming relational products. But we assume that there are only finitely many primitive properties within the ken of the language, at least to begin with.

¹⁵As will soon become clear, these are not to be thought of in the standard model-theoretic way sketched in the first part of this paper.

to monadic predicates. Even here choices arise. First, is there one most general, all inclusive property? From the graph-theoretic point of view: is the structure a tree, or is it rather a forest (an unrooted or multi-rooted tree)? Second, are there cases in which a primitive property is immediately included in more than one primitive property - has more than one immediate ancestor in the structure? (Are there cases of multiple inheritance among the primitive properties?) This is the question, from a graph-theoretic point of view, as to whether the structure is a tree (rooted or not) or an upper semi-lattice (perhaps rootless).

I need not make a decision on these points; such a choice is up to a user of the scheme I am describing, and his/her decision, in turn, depends on the structure of the domain of application and/or his/her conceptualization of that domain. I should, however, be able to show how an account would be given in each case. So I shall begin with the simplest case, that in which the structure is an honest-to-goodness tree.

The ordering relation which generates the algebraic structure is property inclusion (or property entailment), taken as primitive. There is another significant semantic relation: the relation among the immediate descendants of a given node, that is, among siblings. Such properties are taken mutually to exclude one another. (One can then define property independence in terms of inclusion and exclusion.) The intuition, here, is as follows: imagine a portion of the structure which begins with a node for the property of being a mammal. This has, say, 10 immediate descendants marked as mutually exclusive, among which are, e.g., the property of being a cat, the property of being a pig, etc. Each includes the property of being a mammal, and each excludes all the others. Crucially, there is no requirement that

the sequence of mutually exclusive immediate descendants exhausts the immediately superior property. So also, in the case at hand, there is no assumption that the language has primitive predicates for each of the mammalian species. Rather, we want to allow for the discovery of new primitive properties (new species), and their introduction into the language, though not perhaps without limit. I shall return to this point in a moment.¹⁶

Given the above, it's fairly easy to see how the structure of properties should constrain the assignments of extensions to the primitive predicates, those predicates which are associated with nodes in the structure. Such predicates are interpreted, first, by assigning them properties (intensions). The structure among these properties, then, generates the relations among the extensions (sets) associated with the predicates. Specifying a model for such a system would go as follows: first, the entire

¹⁶As should by now have been made clear, and in case it hasn't: "primitive" does not mean "simple". The guiding intuition is that the relation between genus and species is like that between determinables, such as the property of being colored, and determinates, such as the property of being red. The latter includes the former; but the property of being red is not to be analyzed as consisting of the property of being colored together with some other property (unless that other property is that of being red). Nor is the property of being colored to be analyzed as a(n infinite) disjunction of its determinates. Indeed, it does not include any of them; though, to repeat, each of them includes it. In general, more general primitives (determinables) are simpler than the less general primitives (determinates). For a related, but different view, see [34, 35].

domain is assigned to the topmost primitive.¹⁷ The extension assigned to a node D which is an immediate descendant of node A must be a subset of the extension assigned to A. Moreover, given that D belongs to an n-tuple of immediate descendants of A marked as mutually exclusive, the sets assigned to the members of the n-tuple must be disjoint subsets of the set assigned to A. In general, however, this family of sets will not exhaust the set assigned A. (This, of course, is the extensional reflection of the point made above with respect to the intensional structure.)

Suppose, for example, that there are 10 immediate descendants of MAMMAL, marked as mutually exclusive. There will be admissible models in which the extension assigned MAMMAL is not simply the union of the extensions assigned its immediate descendants. The primary fact, however, is that the property of being a mammal is not the "logical (conceptual) sum" of its immediate subordinates. It is not, that is, definable in terms

¹⁷I'm assuming, for the sake of simplicity and familiarity, a standard first-order model structure. In earlier versions of this work, I had not attempted to satisfy the seemingly deep-seated psychological need for specifying such a structure for the language. One alternative was to eschew sets as extensions of predicates altogether and stick with properties; another, a variant of suggestions due to van Fraassen [34, 35], was set-theoretic with a vengeance. That I treat the relationship between the intensional structure and the class of admissible models in the way I do here is due almost entirely to the philosopher and logician George Smith. For more on this, see the acknowledgement at the end.

of them; it is not - ex hypothesi - definable at all.¹⁸

To accommodate multiple inheritance - to allow our structure to be an (upper) semilattice - we need enter only one amendment. Where D is immediately subordinate to A_1, \dots, A_n , the extension assigned to D is, in general, a proper subset of the intersection of the extensions assigned to its immediate ancestors. Again, there can be models in which the extension assigned to D is the intersection of the sets assigned to A_1, \dots, A_n ; but we are assuming that D is a primitive and thus that it is not definable in terms of the other primitives.¹⁹

Finally, we shall take a quick look at the unrooted (or multi-rooted) option. In such a case, should the assignments to the n topmost nodes be independent, identical, or should they rather be mutually exclusive? The first of these options best fits the natural treatment of assignments to parallel and independent nodes. The question to be asked, of course, is what intuitions motivate the move to a forest? My own perusal of the relevant literature convinces me that there is one kernel

¹⁸Note this has the consequence that there can be models in which there are individuals which are assigned to no other predicate than the one correlated with the highest node(s), and so on down for intermediate predicates; e.g. individuals that are - in a given model - mammals, without belonging to any specified species of mammal. For a like-minded view, see [14].

¹⁹We can, of course, introduce a term for the non-primitive property the reader might have had in mind - $(\text{LAMBDA } (x)(A_1(x) \& \dots \& A_n(x)))$; just as in the case discussed above, we could have $(\text{LAMBDA } (x) (A_1(x) \vee \dots \vee A_n(x)))$. Reminder: these lambda expressions are property-designators; they are not to be confused with the complex predicates with which they are correlated. More on complex properties and complex predicates below.

intuition, which yields, at the very least, a forest with two topmost nodes. The kernel intuition is our old friend, intuition (3), in disguise. There is to be one tree for kinds of things and another for qualities of things. Kinds must be distinguished from qualities; being a cat must be distinguished (in kind, no doubt) from being red. (More on this in a moment.)

If this is the operative intuition, I would assume that the assignments would go as follows: the whole domain is assigned to the topmost node in the "thing-hierarchy". Within that hierarchy things go as before. But now what to do with the quality subtree? If one buys the third intuition in a strong form, qualities don't belong to particular things; but only to particular F's or G's, where "F" and "G" are schematic placeholders for common nouns. To buy this intuition, then, is to hold that all predications (attributions) of a quality to an individual are really attributions of the quality to a particular instance of some kind. Put another way: this intuition leads one to deny that kinds of things and qualities of things should be treated as on a par - as both independently determining extensions. (There is such an entity as the set of cats: a set of things which are cats; but there is no such set as the set of things which are red. Still, there is a set of cats which are red.) So, the appropriate move would be to treat the hierarchy of qualities in a quite different way than the taxonomy of kinds: the predicates correlated with these nodes do not get assigned extensions. Such predicates are treated as belonging to a different logical type than the predicates correlated with kind-properties. A plausible line is to treat qualities as functions from kinds to kinds, and in the first instance, as functions from

primitive kinds to defined kinds.²⁰

This last discussion raises the possibility of yet another deviation in syntax from standard first order languages in the direction of sortal quantifiers. We might require that all quantification, indeed all variable binding, including lambda-abstraction, be kind-restricted. Here there are options: one has been explored by Barwise and Cooper [2]; another, closer on the surface to standard first-order notation has been advanced by Gupta [11]. Distinguish between predicates, properly so called, and common nouns, and stipulate that if and only if C is a common noun, x a variable and F a formula, then $((Vx, C) F)$ is a formula. Also, iff as above, then $(\text{LAMBDA } (x, C) F)$ is a complex, defined predicate - correlated with the defined property of being a C which is F . The crucial point is that variable binding requires a sortal. No predication without classification; no quantification without classification

²⁰Syntactically, then, "adjectives", quality-predicates, are of the category CN/CN; this, of course, is the treatment they receive in Montague style grammars for natural languages.

either.²¹

Two last points. First, I have said nothing about whether the formation of complex properties by qualification of simple kinds is unrestricted. For instance, I have said nothing to rule in or out such properties as being a number which is red or being

²¹A few remarks are in order here. It would be misleading to suggest that the only deviations to be encountered here were syntactic. Roughly in order of increasing distance from the familiar: We can opt for sortal quantification, but allow both those predicates associated with sorts or kinds and those associated with qualities to determine sets as extensions. Needless to say, not all qualities are created equal; compare "male criminals" with "alleged criminals" with "tall criminals". Arguably, only the first could plausibly be treated in terms of set-theoretic intersection anyway. For more on this, see [13]. To remain as close as possible to standard first-order model structures, we would then simply posit that the intensions associated with the first kind of monadic predicate were a special kind of property, viz. that of belonging to a particular (natural) kind, leaving this last notion primitive so far as the specification of the semantic account is concerned. This is what we have done with the notion of a property; we have most emphatically not explicated this notion in the standard modal model-theoretic way as a function from possible worlds to sets of (possible) individuals. Just as in standard model-theoretic treatments of modality the notion of possible world is unexplicated. The next option should now be obvious; we can move to a modal model structure and assign different semantic types to predicates for kinds and predicates for qualities. For a version of this move, less radically deviant than Bressan's, see Gupta's "The Logic of Common Nouns". More radical still is the analysis in terms of function application. One alternative would be a variant of the account given by Barwise and Cooper; this has the attractive feature of not going unnecessarily modal. Here, too, we would take the notion of (natural) kinds as primitive from the point of view of semantic theory.

an idea which is green.²² This point is connected with the second one. Once we have split up our tree into a forest, why not split up our sub-trees. So, for instance, why assume a single topmost node in the kind hierarchy, a node, presumably, for the property of being a thing? Surely, we can split here as well, separating off a tree of physical things from a tree of abstract entities, etc. We can, then, do the same for qualities and somehow express selection restrictions blocking the anomalous "kinds" instanced above. All these things are possible, and some of them may even be desirable. Here, too, I shall beg off pursuing the details of the various options open to us.

I won't now go into any details of the ways of extending our structures to handle the correlates of many-placed predicates. I do, however, assume that there are only finitely many primitives, so for some n , there will be no primitive relations of degree n or greater. Relations of such degrees there will be, generated (e.g.) by taking relational products of relations of lesser degree. There are, of course, a number of questions to be answered and a number of options to be explored. E.g., shall we allow inclusion relations between relations of different degree; and if so, how shall we handle such cases? How shall we mark, e.g., the fact that for a given ordered pair $\langle x, y \rangle$ to be in the extension associated with a certain two-place predicate, the first element must have a certain property or belong to some

²²From a syntactic point of view: I have not discussed restrictions on that mode of complex "predicate" forming which consists of applying an adjectival modifier to a kind term. Semantically speaking, I'm avoiding giving a semantical account of sortal incorrectness or anomaly. For one such account, see [33].

specified kind? I will simply note here that there are no significant obstacles to the extension to polyadic relations, whichever of the options available to us in the monadic case we decide to extend to the polyadic.

Finally, a brief mention of complex, defined properties and their associated complex predicates. From what was said above, it should be clear that there are options as well with respect to the kinds of complex concept forming operators we allow. For instance, if one opts for the single rooted structure, one may opt for standard, unrestricted Boolean compounds. So, one would expect a node for the defined property of being a red cat: (LAMBDA (x) (cat(x) & red(x))), which in every admissible model will be assigned the intersection of the sets assigned to its two constituent primitives. If, on the other hand, one is attracted to the picture suggested by intuition (3), the relevant complex concept-forming functor will be of the CN/CN variety, that is, the compound will really be a result of function application. Moreover, one might want to block Boolean compounds of (primitive) kind terms (e.g. both (LAMBDA (x) (cat(x) & dog(x))) and (LAMBDA (x) (flower(x) v mouse(x))). Surely, one will want to allow some such combinations of qualities, even if only sortally restricted combinations; e.g. (LAMBDA (x, cat) (white(x) v black(x))). But, then what of (LAMBDA (x, cat) (white(x) & black(x))) - assuming WHITE and BLACK are marked as mutually exclusive under COLORED? Once all these decisions are made, one can specify the nature of the resulting structure of all possible properties - primitive and defined. Once again, I forego the messy details.

Enough has been said, I hope, to make fruitful an attempt to answer the question that will naturally arise in the mind of anyone used to a more orthodox account. Why not capture all

these intuitions about semantic interdependencies among primitives by way of meaning postulates? This last is the traditional alternative mode of capturing the intuition that items in the extralogical vocabulary can partially determine the implication relations among sentences.²³

There are, I think, reasons for preferring the semantic network picture. First, the meaning postulate account denies (or does not account for) the distinction between analytic sentences, whose necessity is grounded in the internal structure of their non-logical vocabulary, and either logical truths or non-logical necessary truths. So for instance, the account takes no official cognizance of the distinction between sentences such as "Every even number is the sum of two primes", and "All cats are mammals". The first, if true, is a necessary truth about the natural numbers; it is not a logical truth. The second, according to a particular instance of the account I have sketched, is rendered true solely by the full specification of the language to which that sentence belongs; yet it too is not, in the standard sense, a truth of logic. Nor does the traditional account distinguish either of the above two sentences from "All squares are rectangles" or "All squares are either

²³The idea is due to Carnap and has been extensively exploited by Montague and his followers. See [7, 22]. Ron Brachman and myself raised this possibility in [12]. See also [14].

squares or red".²⁴

To return for a moment to pure description. If someone wanted to specify as part of the determination of his/her language that the predicate associated with the property of being a mammal was analytically entailed by that associated with the property of being a cat, one could do so by arranging the intensional structure associated with the language in the right way. And this dependency between the two primitive predicates would be something over and above the necessity of the associated universalized conditional sentence.²⁵

The first reason for preferring a semantic network account

²⁴The differences alluded to above are tricky matters. Why not claim that the necessary truth of the second sentence is grounded in facts about the world of living things and their essences? This is Kripke's view in [17]. I am very sympathetic to this rejoinder; still, at least one point can be made on my behalf. I am simply presenting for the reader's delectation and delight a view according to which there is a distinction between sentences whose necessity is grounded in extra-linguistic matters of fact and those (perhaps also) grounded in the very structure of the language itself; though not in the way the truths of logic are so grounded. According to the present proposal, one can accommodate these distinctions. So much by way of "argumentation".

²⁵The arcs in the conceptual structure between super- and subordinate nodes do not realize the complex syntactic operation of forming the universalized conditional sentence in whose matrix the open sentence whose predicate is that correlated with the subordinate node is the antecedent and whose consequent is the open sentence associated with the superordinate node. Nodes and links are one thing; sentences another. Of course, the abstract syntax of the language could be realized by a set of graph-type structures and operations thereon, rather than the usual linear, one-dimensional structures. To take this last point as important is to trivialize the first of our three guiding intuitions.

to one in terms of meaning postulates introduces a second. One major difference between the non-standard account sketched and the traditional alternative is that meaning postulates are simply sentences of the language picked out from their siblings in just the same way that axioms are picked out in standard axiomatizations - namely, by being stuck in a list with the metalinguistic heading "AXIOM" ("MEANING POSTULATE") prefixed to it. The crucial point here is that, in the standard accounts, one uses the meaning postulates precisely as axioms. The set of "analytic" sentences is identified with the set of (standard) logical consequences of the postulates; for the finicky, the purely logical truths are often subtracted. Thus the meaning postulates enter into the explication of analyticity solely in virtue of their logical form plus their tag. The interpretations of the syntactic primitives covered by the postulates are not incorporated into the specification of the language. In the nonstandard account, on the other hand, the primary necessities are not formal; the relation of "predicate inclusion" that holds between (e.g.) "cat" and "mammal" is not a function of their logical forms; they are both syntactically simple predicates and hence are devoid of significant logical form.²⁶ In this respect, compare " $(x) (Crow(x) \rightarrow Bird(x))$ " and " $(x) (Crow(x) \rightarrow Black(x))$ ". These sentences are of exactly the same syntactic type; but, again on at least one plausible instance of the scheme sketched, the truth of the first is grounded in property inclusion and, thereby, in the language itself; not so that of the second.

²⁶I trust the reader will pardon the sloppiness involved in talk of "predicate" as opposed to "property" inclusion.

This second point leads us to the third and penultimate point. The relations of inclusion or entailment between primitive predicates are not, we have noted, a matter of logic at all. Rather, they are determined solely by the structure of primitives; and, in principal, all such relations can be "read right off" from the structure. Indeed, a central feature of the account given here is that it allows us to make a sharp distinction between matters of logic and matters of language. Matters of logic are matters of logical form; there is no room for such concerns to intrude on the inclusion relations among the primitive predicates. Computationally, matters of logic relate to inferential operations; more concretely, to the application of rules of transformation. These last are aptly named. Where there is deduction, there are steps, steps which involve the manipulation and generation (broadly speaking, the transformation) of terms or of formulae. There are no such steps involved with respect to determining the crucial semantic relationships among primitive predicates; there is nothing beyond lookup. That is, we stipulate that the algebraic structure of primitives must be specified completely in advance. That specification can be seen as being realized in a number of ways; but in the abstract it is as if attached to each node was a specification of the paths on which it is located, and of course, its relative location on that path. There is no computation to be done beyond seeing whether a given other predicate does or not appear on such a path (and where). If this latter does, in a particular implementation, require computation, that is solely a function of the implementation. To put the point in yet another, and clearer, way: the computational steps involved in actually traversing the tree and establishing relative locations among a given set of nodes are not to be identified with the computational steps involved in implementing a proof procedure

which is sensitive to the logical forms of the expressions over which its operations are defined.

Finally, one minor point about the alternative account using meaning postulates. Let's look at an extremely simple case in which there is one node, A , which immediately dominates two others, A_1 and A_2 . We assume that these two subordinate nodes are marked as mutually exclusive. What would the set of meaning postulates for these primitives be? At first, things are simple: " $(x) (A_1(x) \rightarrow A(x))$ ", " $(x) (A_2(x) \rightarrow A(x))$ ", " $(x) (A_1(x) \rightarrow \neg A_2(x))$ ". Now, how about " $\neg (x) (A(x) \rightarrow (A_1(x) \vee A_2(x)))$ "? By itself, this won't do. Though it's crucial that " A "'s extension not be determined in all models to be simply the union of the extensions assigned to " A_1 " and " A_2 ", we don't want to rule out that the set of A 's actually is, as it happens, just that union. Another option is to take advantage of the resources of set theory. In every case like the one at hand, we can simply introduce a predicate, " A_3 " - which will, of course, be a syntactically simple predicate, for in the standard story, there are no others - and assign to it the following set: $\{x | Ax \ \& \ \neg A_1x \ \& \ \neg A_2x\}$. The set assigned to " A_3 " can be empty or not. This assignment goes somewhat against the grain of the standard story, since it is completely a function of the assignments to the other predicates, and because now the same is true of the assignment to " A ".

Moreover, note the following: if we stipulate that predicates are, in the first place, correlated with properties, and if we take seriously the distinction between primitive and defined properties, then there is no reason to assume that there is a primitive predicate (one correlated with a primitive property) whose associated extension just is the logical difference between " A "'s extension and the union of the

extensions of "A1" and "A2". Indeed, the only way we might have of representing the extension of "A3" is by way of a set abstract in which the subsuming predicate ("A") makes an essential appearance. (Surely we can assume that we can't always present the relevant sets by lists of their members.)²⁷

There are other options, of course. One is to go modal. First, we prefix all of the foregoing (except the denial of exhaustion of A's) with a necessity operator. Then, what we want is that there can be A's which are neither A1's nor A2's. I.e. "It is not necessary that $(x) (A(x) \rightarrow (A1(x) \vee A2(x)))$ ". (Or: "It is possible that there be an A which is neither an A1 nor an A2".) The modal move is not unnatural; but it strikes me as better not to introduce modalities until one has to, to handle explicit modal operators for example. Let's not muck about with other worlds until we've done what we can with the resources available to us in this one.

I shall refrain from going on at any length at all about the most appropriate account of the logical connectives in the kind of scheme I've sketched. In a sense, what I've said about the non-logical vocabulary can be seen as imposing no constraints on a semantic account of the logical vocabulary, especially if we confine our attention to logical operations on sentences, and ignore the use of logical operators in forming complex, defined

²⁷I have presumed that it is a condition of adequacy that the primitive character of the properties be captured by the postulates. This requirement, which is precisely what causes trouble for the account in terms of meaning postulates, might be challenged.

predicates.²⁸

²⁸I must 'fess up, though; I think there may well be more or less appropriate accounts of the connectives in this regard. In particular, if one wants to take the inclusion metaphor seriously, a natural family of logics to look into are the Relevance Logics of the Pittsburgh School of Logic. See [1, 34].

Bolt Beranek and Newman Inc.

Report No. 5117

4. TECHNICAL ADDENDUM: TAXONOMIES OF NATURAL KINDS

Talk of taxonomic structures can be taken in various ways. I have aimed to take it seriously. To that end, I want here to look more closely at the structure of taxonomies all of whose nodes are natural kinds (or the property of belonging to such-and-such a natural kind). Professor John Mylopoulos of Toronto was kind enough to point me in the direction of research done by biologists, zoologists and anthropologists on the theory of such taxonomies. He did this by referring me to, and then giving me a copy of, a paper by the Berkeley anthropologist Paul Kay [14].

Kay directs us as follows:

Consider the graph-theoretic structure in which taxa (the elements of taxonomies - D.I.) are represented by nodes and immediate precedence by directed edges. In all previous treatments of formal taxonomic structure of which I am aware, the resulting graph is a singly-rooted (and directed) tree. That is, there is exactly one vertex with no incoming arrows and no vertex with more than one incoming arrow. If the latter condition did not hold, we could get subgraphs containing cycles..., and by definition a tree has no cycles. [14] p. 157.

Kay, however, wants to allow such "cycles" in taxonomic structures and I want to join him in this.²⁹ His arguments for dropping the restriction to trees allude to actual data, culled from the field, in the form of responses to questions designed to

²⁹As Kay remarks: "There almost certainly are empirical limitations on the kinds of cycles that may occur in taxonomic graph, but an investigation of that topic would not be in order here." [14]. Or here. By the way, I take it that it's clear that Kay has in mind lattice-type structures, and not "cyclic" or "circular" graph-structures.

elicit the "folk taxonomies" of real, live folk. They allude, as well, to the proscribed (by me, for now) notion of prototypicality. My motivation has been to account for what in the AI literature has come to be called "multiple inheritance". The resulting structures are (join or) upper semilattices. That is, they are triples $\langle S, R, T \rangle$, where S is a set, R a reflexive, transitive and antisymmetric relation on S (usually pronounced "less than or equal to"), and T an element of S such that for every s in S , sRT . Moreover, for every two elements s, s' of S , there is a unique element of S (designated their join or least upper bound: $s \cup s'$) which is the least element s'' such that both sRs'' and $s'R s''$.³⁰

Once having gone this far, we may want to go one step farther and stipulate that all such structures will have a least element \exists (or BOTTOM) such that for all s in S , $\exists Rs$. This ploy, which can be seen as just that - a technical ploy - has the nice effect of turning our structures into lattices. We now have it that for every pair of elements of S , s and s' , there is a unique element of S (designated as their meet or greatest lower bound: $s \wedge s'$) which is the greatest element s'' in S such that both $s''Rs$ and $s''Rs'$.³¹

³⁰The classic work on lattice theory is [3]. I should note that we impose the condition that no node has two immediate predecessors of the same rank, where this term is defined as in [14], ps. 160f.

³¹I should note that $T(OP)$ can also be looked on as a technical ploy; that is, we need not be committed to the view that there is a highest genus, a most general natural kind. We can treat both $T(OP)$ and BOTTOM as "ideal elements" of the lattice, thrown in to the structure of natural kinds to give it a more pleasing and tractable mathematical shape.

Not all lattices are created equal. Are there further, natural conditions which taxonomic lattices must meet? There are at least two - both cited by Kay and intuitively appealing on their faces.³² Both are best illustrated before being formally specified. Consider the simple structure that consists of the node for mammal and n immediate descendants, one for each of the n known species of mammal. First, nothing can be both a dog and a whale; and in general, the immediate descendant species under a genus (or higher species) are thought to be mutually exclusive. Second, there is no presumption that every mammal must belong to one of the n species; that is, there is no presumption that the immediate descendants under a genus (higher species) induce a partition of their ancestor. (See [14], ps. 155f.)

What is the effect of these conditions (especially the first) on the characteristics of taxonomic lattices. Well, let us first formally specify (and generalize) the condition of mutual exclusiveness:

for any two elements of s, s' of S , if s and s' are incomparable (if neither sRs' nor $s'R s$), then no thing can be characterized by their greatest lower bound.

More formally:

for any two elements s, s' of S , either $s \wedge s' = s$ or $s \wedge s' = s'$ or $s \wedge s' = \text{BOTTOM}$.

With this in hand, we can ask, first, if taxonomic lattices

³²If I may say so, I imposed both of them in the work I had done prior to reading Kay's piece.

are distributive.³³ Surely not: $\text{homo sapiens} \cup (\text{frog} \wedge \text{pig}) = \text{homo sapiens}$; while $(\text{homo sapiens} \cup \text{frog}) \wedge (\text{homo sapiens} \cup \text{pig}) = \text{animal}(?) \wedge \text{mammal} = \text{mammal}$. (My thanks to the Muppets for this and other examples.) Are taxonomic lattices at least modular? A lattice is modular iff

$$\text{for all } a, b, c: \text{ if } aRc, \text{ then } a \cup (b \wedge c) = (a \cup b) \wedge c.$$

No, I am afraid not (and Kay, in the diagrams on the top of p. 158, comes within one well-placed node of giving us an example of a non-modular lattice). Technicalities aside, consider the following example: homo sapiens is less than or equal to mammal , but $\text{homo sapiens} \cup (\text{frog} \wedge \text{mammal}) = \text{homo sapiens}$ which does not equal $(\text{homo sapiens} \cup \text{frog}) \wedge \text{mammal} = \text{mammal}$.³⁴

Needless to say, the lattice is not complemented. Indeed, (taxonomies of) natural kinds are (nowhere near) closed under no logical operations. Informally, if we grant ourselves property forming operators analogous to the truth-functional operators of our language, then if A and B are natural kinds (construed if you like in terms of properties; i.e., the property of being an A, the property of being a B) - in general - neither the property of

³³A lattice is distributive just in case for any any a, b, c , $a \cup (b \wedge c) = (a \cup b) \wedge (a \cup c)$ and $a \wedge (b \cup c) = (a \wedge b) \cup (a \wedge c)$.

³⁴This latter result is unhappy, simply because modular lattices are such ubiquitous critters in algebra, and hence are terribly well-studied. Taxonomic lattices do, however, have the following property: if x is strictly between **BOTTOM** and a and a and b are incomparable, then $x \cup b = a \cup b$. This, it seems to me, is a nice property for something called a hierarchy to have.

not being an A nor that of not being a B is a natural kind property; nor is that of being either an A or a B, nor that of being both an A and B. One should not confuse the lattice operations of meet and join on a taxonomic lattice either with the truth-functional (logical) operations of conjunction and disjunction or with the set-theoretic operations of intersection and union.

Two last points: first, taxonomic structures need not be finite; but it is surely all right to assume that those that are taxonomies of natural kinds are finite. Second, there may be perfectly good uses for taxonomic structures defined over many-placed relations, even if there are no relational natural kinds.

Bolt Beranek and Newman Inc.

Report No. 5117

5. ACKNOWLEDGEMENTS

Work on this paper was inspired by discussions, over 3 years, with Rusty Bobrow, Ron Brachman, Hector Levesque, Jim Schmolze, Brian Smith, Bill Woods and (other) members of what might be called the KL-ONE WORKING GROUP. It is from them that I learned whatever I know about semantic networks. Thanks also to John Mylopoulos and other members of the Computer Science Department at the University of Toronto who were subjected to this material in not yet digested form and let me live to tell about it; and to Bob Moore for a stimulating discussion of the material.

With respect to the section on inheritance, a special acknowledgement is due to George Smith, of the Tufts University Philosophy Department. It was he who convinced me of the utility of connecting up the account in terms of property inclusion with an account of the constraints on admissible models of the standard first-order variety. It was he, as well, who convinced me that the connection between property-inclusion and analytic entailments a la Anderson-Belnap really could be given a natural and plausible theoretical grounding for the case of natural languages. The present work (which, as noted, touches neither on purely logical matters nor on issues in the semantics of natural languages) would have been even more deviant in conception than it is, if he hadn't shared with me his current work (in the form of chapters of a manuscript written in collaboration with Jerrold Katz) on the semantics of natural languages (see [32]).

REFERENCES

- [1] Anderson, A. and Belnap, N.
Entailment.
Princeton University Press, Princeton, 1975.
- [2] Barwise, J. and Cooper, R.
Generalized Quantifiers and Natural Languages.
Unpublished manuscript.
- [3] Birkhoff, G.
AMS Colloquium Publications. Volume XXV: Lattice Theory.
American Mathematical Society, New York, 1948.
revised 2nd edition.
- [4] Brachman, R. J.
A Structural Paradigm for Representing Knowledge.
Technical Report 3605, Bolt Beranek and Newman Inc., May,
1978.
- [5] Brachman, R. J.
An Introduction to KL-ONE.
In Brachman, R. J., et al. (editors), Research in Natural
Language Understanding, Annual Report (1 Sept. 78 - 31
Aug. 79), pages 13-46. Bolt Beranek and Newman Inc,
Cambridge, MA, 1980.
- [6] Bressan, A.
A General Interpreted Modal Calculus.
Yale University Press, New Haven, 1972.
- [7] Carnap, R.
Meaning and Necessity.
University of Chicago Press, Chicago, 1955 (2nd ed.).
- [8] Church, A.
On Carnap's Analysis of Statements of Assertion and Belief.
Analysis X:97-99, 1950.
- [9] Church, A.
Intensional Isomorphism and Identity of Belief.
Philosophical Studies V:65-73, 1954.
- [10] Cresswell, M.
Hyperintensional Logic.
Studia Logica 34(1):25-38, 1975.

- [11] Gupta, A.
The Logic of Common Nouns.
Yale University Press, New Haven, 1980.
- [12] Israel, D. and Brachman, R.
Distinctions and Confusions: A Catalogue Raisonne.
In Proceedings of the Seventh International Joint
Conference on Artificial Intelligence, pages 452-458.
IJCAI, Vancouver, 1981.
- [13] Kamp, J. A. W.
Two Theories about Adjectives.
In Keenan, E. L. (editor), Formal Semantics of Natural
Language, pages 123-155. Cambridge University Press,
Cambridge, 1975.
- [14] Kay, P.
A Model-Theoretic Approach to Folk Taxonomy.
Social Science Information 14(5):151-166, 1975.
- [15] Kripke, S.
A Completeness Theorem in Modal Logic.
Journal of Symbolic Logic 24:1-14, 1959.
- [16] Kripke, S.
Semantical Analysis of Modal Logic I.
Zeitschrift fur Mathematische Logik und Grundlagen der
Mathematik 9:67-96, 1963.
- [17] Kripke, S.
Naming and Necessity.
In Davidson, D. and Harman, G. (editors), Semantics of
Natural Language, pages 253-355. D. Reidel, Dordrecht-
Holland, 1972.
- [18] Mates, B.
Synonymy.
In Linsky, L. (editor), Semantics and The Philosophy of
Language, pages 111-136. University of Illinois Press,
Urbana, 1952.
- [19] McCarthy, J.
First Order Theories of Individual Concepts and
Propositions.
In Hayes, P.J., Michie, D., and Mikulich, L. (editors),
Machine Intelligence 2, pages 129-147. Halsted Press,
New York, 1979.

- [20] Montague, R.
Logical Necessity, Physical Necessity, Ethics, and
Quantifiers.
In Thomason, R. (editor), Formal Philosophy, pages 71-83.
Yale University Press, New Haven, 1974.
- [21] Montague, R.
Pragmatics and Intensional Logic.
In Thomason, R. (editor), Formal Philosophy, pages 119-147.
Yale University Press, New Haven, 1974.
- [22] Montague, R.
The Proper Treatment of Quantification in Ordinary English.
In Thomason, R. (editor), Formal Philosophy, pages 247-270.
Yale University Press, New Haven, 1974.
- [23] Moore, R. C.
Reasoning about Knowledge and Action.
Technical Report 191, Artificial Intelligence Center, SRI
International, October, 1980.
- [24] Reiter, R., and Criscuolo, G.
Some Representational Issues in Default Reasoning.
Technical Report 80-7, Dept. of Computer Science, The
University of British Columbia, August, 1980.
- [25] Schubert, L. K., Goebel, R. G., and Cercone, N. J.
The Structure and Organization of a Semantic Net for
Comprehension and Inference.
In Findler, N. V. (editor), Associative Networks:
Representation and Use of Knowledge by Computers, pages
121-175. Academic Press, New York, 1979.
- [26] Scott D.
Advice on Modal Logic.
In Lambert, K. (editor), Philosophical Problems in Logic,
pages 143-172. D. Reidel, Dordrecht-Holland, 1970.
- [27] Shapiro, S.C.
A net structure for semantic information storage,
deduction, and retrieval.
In Proceedings of the Second International Joint Conference
on Artificial Intelligence, pages 512-523. IJCAI,
London, England, 1971.
- [28] Shapiro, S.C.
The SNePS Semantic Network Processing System.
In Findler, N.V. (editor), Associative Networks:

Representation and Use of Knowledge by Computers, pages 179-203. Academic Press, New York, 1979.

- [29] Shapiro, S.C.
Using Non-Standard Connectives and Quantifiers for
Representing Deduction Rules in a Semantic Network.
Presented at: Current Aspects of AI Research, a seminar
held at the Electrotechnical Laboratory, Tokyo.
- [30] Shapiro, S.C.
What do Semantic Network Nodes Represent?.
Technical Report 7, Department of Computer Science, State
University of New York at Buffalo, July, 1981.
- [31] Shapiro, S.C., and Maida, A.S.
Intensional Concepts in Propositional Semantic Networks.
Technical Report 171, Department of Computer Science, State
University of New York at Buffalo, February, 1981.
- [32] Smith, G.E. and Katz, J.J.
Intentionally Admissible Models: The Extensional
Interpretation of Intensional Semantics.
manuscript in preparation.
- [33] Thomason, R.
A Semantic Theory of Sortal Incorrectness.
Journal of Philosophical Logic 1:209-258, 1972.
- [34] van Fraassen, B.
Facts and Tautological Entailments.
Journal of Philosophy(66):477-487, 1969.
- [35] van Fraassen, B.
Extension, Intention, and Comprehension.
In Munitz, M.K. (editor), Logic and Ontology, pages 101-
131. New York University Press, New York, 1973.

Official Distribution List

Contract N00014-77-C-0371

	<u>Copies</u>
Defense Documentation Center Cameron Station Alexandria, VA 22314	12
Office of Naval Research Information Systems Program, Code 437 Arlington, VA 22217	2
Office of Naval Research Code 200 Arlington, VA 22217	1
Office of Naval Research Code 455 Arlington, VA 22217	1
Office of Naval Research Code 458 Arlington, VA 22217	1
Office of Naval Research Branch Office, Boston 495 Summer Street Boston, MA 02210	1
Office of Naval Research Branch Office, Chicago 536 South Clark Street Chicago, IL 60605	1
Office of Naval Research Branch Office, Pasadena 1030 East Green Street Pasadena, CA 91106	1
Dr. Martin Epstein National Library of Medicine Bldg. 38A, 8th Floor Lab. 8600 Rockville Pike Bethesda, MD 20209	1
Naval Research Laboratory Technical Information Division Code 2627 Washington, DC 20380	6

Naval Ocean Systems Center 1
Advanced Software Technology Division
Code 5200
San Diego, CA 02152

Dr. A.L. Slafkosky 1
Scientific Advisor
Commandant of the Marine Corps (Code RD-1)
Washington, D.C. 20380

Mr. E. H. Gleissner 1
Naval Ship Research & Development Center
Computation & Mathematics Dept.
Bethesda, MD 20084

Captain Grace M. Hopper, USNR 1
Naval Data Automation Command
Code OOH - Washington Navy Yard
Washington, DC 20374

Mr. Kin B. Thompson 1
NAVDAC 33
Washington Navy Yard
Washington, D.C. 20374

Advanced Research Projects Agency 1
Information Processing Techniques
1400 Wilson Boulevard
Arlington, VA 22209

Captain Richard L. Martin, USN 1
3311 W. Avenue
Newport News, VA 23607

Director, National Security Agency 1
Attn: R54: Mr. Page
Fort G.G. Meade, MD 20755

Director, National Security Agency 1
Attn: R53, Mr. Glick
Fort G.G. Meade, MD 20755

Major James R. Kreer, Chief 1
Information Sciences
Dept. of the Air Force
Air Force Office of Scientific Research
European Office of Aerospace Research
and Development
Box 14
FPO New York 09510

Mr. Fred M. Griffiee 1
Technical Advisor C³ Division
Marine Corps Development
and Education Command
Quantico, VA 22134

END

DATE
FILMED

11-82

DTIC